

Tagging like Humans: Diverse and Distinct Image Annotation

Baoyuan Wu¹, Weidong Chen¹, Peng Sun¹, Wei Liu¹, Bernard Ghanem², and Siwei Lyu³

¹Tencent AI Lab ²KAUST ³University at Albany, SUNY

wubaoyuan1987@gmail.com, powerchen@tencent.com, pengsun000@gmail.com, wliu@ee.columbia.edu,
bernard.ghanem@kaust.edu.sa, slyu@albany.edu

Abstract

In this work we propose a new automatic image annotation model, dubbed **diverse and distinct image annotation (D²IA)**. The generative model D²IA is inspired by the ensemble of human annotations, which create semantically relevant, yet distinct and diverse tags. In D²IA, we generate a relevant and distinct tag subset, in which the tags are relevant to the image contents and semantically distinct to each other, using sequential sampling from a determinantal point process (DPP) model. Multiple such tag subsets that cover diverse semantic aspects or diverse semantic levels of the image contents are generated by randomly perturbing the DPP sampling process. We leverage a generative adversarial network (GAN) model to train D²IA. Extensive experiments including quantitative and qualitative comparisons, as well as human subject studies, on two benchmark datasets demonstrate that the proposed model can produce more diverse and distinct tags than the state-of-the-arts.

1. Introduction

Image annotation is one of the fundamental tasks of computer vision with many applications in image retrieval, caption generation and visual recognition. Given an input image, an image annotator outputs a set of keywords (tags) that are relevant to the content of the image. Albeit an impressive progress has been made by current image annotation algorithms, to date, most of them [31, 26, 12] focus on the *relevancy* of the obtained tags to the image with little consideration to their inter-dependencies. As a result, algorithmically generated tags for an image are relevant but at the same time less informative, with redundancy among the obtained tags, *e.g.*, one state-of-the-art image annotation algorithm ML-MG [26] generates tautology ‘*people*’ and ‘*person*’ for the image in Fig. 1(f).

This is different from how human annotators work. We

illustrate this using an annotation task involving three human annotators (identified as A1, A2 and A3). Each annotator was asked to independently annotate the first 1,000 test images in the IAPRTC-12 dataset [8] with the requirement of “describing the main contents of one image using as few tags as possible”. One example of the annotation results is presented in Fig. 1. Note that individual human annotators tend to use semantically *distinct* tags (see Fig. 1 (b)-(d)), and the semantic redundancy among tags is lower than that among the tags generated by the annotation algorithm ML-MG [26] (see Fig. 1(f)). Improving the semantic distinctiveness of generated tags has been studied in recent work [23], which uses a determinantal point process (DPP) model [10] to produce tags with less semantic redundancies. The annotation result of running this algorithm on the example image is shown in Fig. 1(g).

However, such results still lack in one aspect when comparing with the annotations from the ensemble of human annotators (see Fig. 1(e)). The collective annotations from human annotators also tend to be *diverse*, consisting of tags that cover more semantic elements of the image. For instance, different human annotators tend to use tags across different abstract levels, such as ‘*church*’ vs. ‘*building*’, to describe the image. Furthermore, different human annotators usually focus on different parts or elements of the image. For example, A1 describes the scene as ‘*square*’, A2 notices the ‘*yellow*’ color of the building, while A3 finds the ‘*camera*’ worn on the chest of people.

In this work, we propose a novel image annotation model, namely diverse and distinct image annotation (D²IA), which aims to improve the diversity and distinctiveness of the tags for an image by learning a generative model of tags from multiple human annotators. The distinctiveness enforces the semantic redundancy among the tags in the same subset to be small, while the diversity encourages different tag subsets to cover different aspects or different semantic levels of the image contents. Specifically, this generative model first maps the concatenation of the image

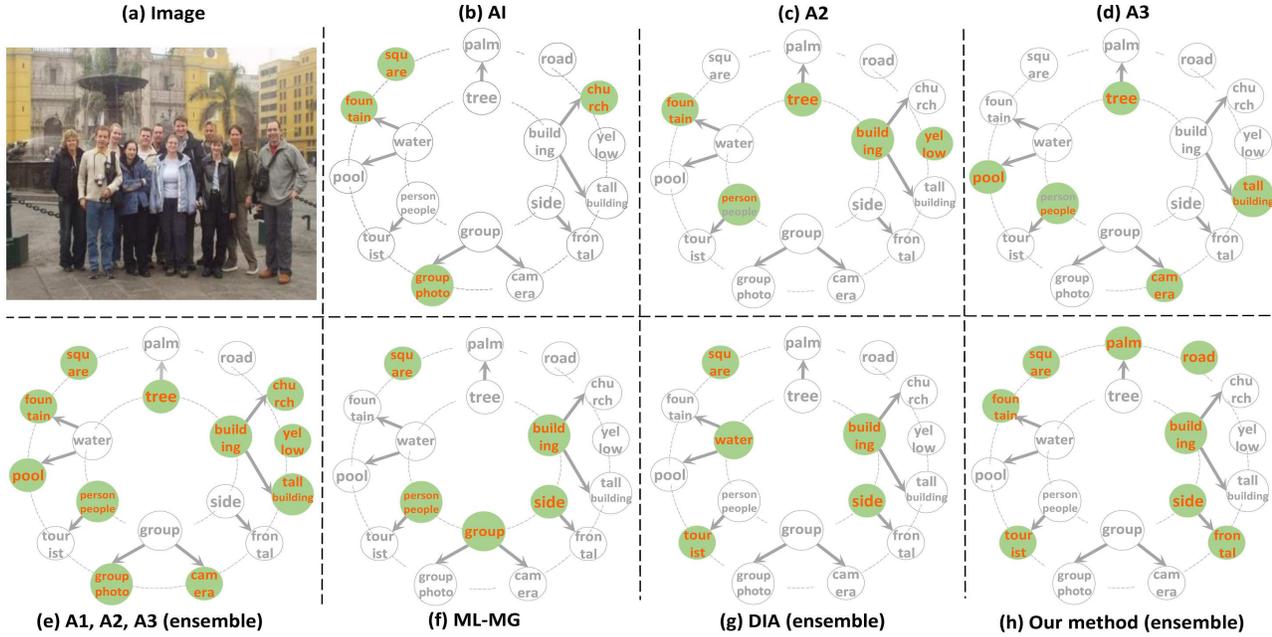


Figure 1. **An example illustrating the diversity and distinctiveness in image annotation.** The image (a) is from IAPRTC-12 [8]. We present the tagging results from 3 independent human annotators (v)-(d), identified as A1, A2, A3, respectively, as well as their ensemble result (e). We also present the results of some automatic annotation methods. ML-MG [26] (f) is a standard annotation method that requires the relevant tags. DIA (ensemble) [23] (g) indicates that we repeat the sampling of DIA for 3 times, with the requirement that each subset includes at most 5 tags, and then combine these 3 subsets to one ensemble subset. Similarly, we obtain the ensemble subset of our method (h). In each graph, nodes are candidate tags and the arrows connect parent and child tags in the semantic hierarchy. This figure is better viewed in color.

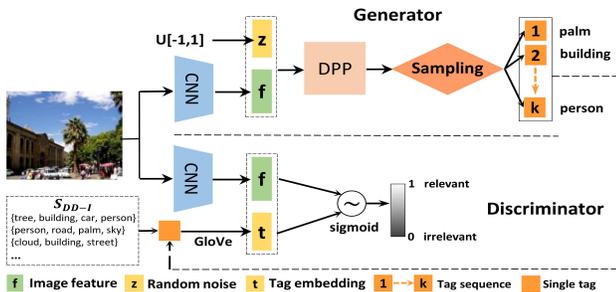


Figure 2. A schematic illustration of the structure of the proposed D^2IA -GAN model. S_{DD-I} indicates the ground-truth set of diverse and distinct tag subsets for the image I , which will be defined in the Section 3.

feature vector and a random noise vector to a posterior probability with respect to all candidate tags, and then incorporates it into a determinantal point process (DPP) model [10] to generate a distinct tag subset by sequential sampling. Utilizing multiple random noise vectors for the same image, multiple diverse tag subsets are sampled.

We train D^2IA as the generator in a generative adversarial network (GAN) model [7] given a large amount of human annotation data, which is subsequently referred to as D^2IA -GAN. The discriminator of D^2IA -GAN is a neural network measuring the relevance between the image feature and the tag subset that aims to distinguish the generated tag subsets and the ground-truth tag subsets from human anno-

tators. The general structure of D^2IA -GAN model is shown in Fig. 2. The proposed D^2IA -GAN is trained by alternative optimization of the generator and discriminator while fixing the other until convergence.

One characteristic of the D^2IA -GAN model is that its generator includes a sampling step which is not easy to optimize directly using gradient based optimization methods. Inspired by reinforcement learning algorithms, we develop a method based on the *policy gradient* (PG) algorithm, where we model the discrete sampling with a differentiable policy function (a neural network), and devise a *reward* to encourage the generated tag subset to match the image content as close as possible. Incorporating the policy gradient algorithm in the training of D^2IA -GAN, we can effectively obtain the generative model for tags conditioned on the image. As shown in Fig. 1(h), using the trained generator of D^2IA -GAN can produce diverse and distinct tags that are closer to those generated from the ensemble of multiple human annotators (Fig. 1(e)).

The main contributions of this work are four-fold. **(1)** We develop a new image annotation method, namely diverse and distinct image annotator (D^2IA), to create relevant, yet distinct and diverse annotations for an image, which are more similar to tags provided by different human annotators for the same image; **(2)** we formulate the problem as learning a probabilistic generative model of tags

conditioned on the image content, which exploits a DPP model to ensure distinctiveness and conducts random perturbations to improve diversity of the generated tags; (3) the generative model is adversarially trained using a specially designed GAN model that we term as D^2IA -GAN; (4) in the training of D^2IA -GAN we use the policy gradient algorithm to handle the discrete sampling process in the generative model. We perform experimental evaluations on ESP Game [20] and IAPRTC-12 [8] image annotation datasets, and subject studies based on human annotators for the quality of the generated tags. The evaluation results show that the tag set produced by D^2IA -GAN is more diverse and distinct when comparing with those generated by the state-of-the-art methods.

2. Related Work

Existing image annotation methods fall into two general categories: they either generate all tags simultaneously using multi-label learning, or predict tags sequentially using sequence generation. The majority of existing image annotation methods are in the first category. They mainly differ in designing different loss functions or exploring different class dependencies. Typical loss functions include square loss [31, 22], ranking loss [6, 12], cross-entropy loss [33], *etc.* Commonly used class dependencies include class co-occurrence [25, 28, 13], mutual exclusion [2, 29], class cardinality [27], sparse and low rank [24], and semantic hierarchy [26]. Besides, some multi-label learning methods consider different learning settings, such as multi-label learning with missing labels [25, 28], label propagation in semi-supervised learning [15, 5, 4] and transfer learning [18] settings. A thorough review of multi-label learning based image annotation methods can be found in [32].

Our method falls into the second category, which generates tags in a sequential manner. This can better employ the inter-dependencies of the tags. Many methods in this category are built on sequential models, such as recurrent neural networks (RNNs), which work in coordination with convolutional neural networks (CNNs) to exploit their representation power for images. The main difference of these works lies in designing an interface between CNN and RNN. In [9], features extracted by a CNN model were used as the hidden states of a RNN. In [21], the CNN features were integrated with the output of a RNN. In [14], the predictions of a CNN were used as the hidden states of a RNN, and the ground-truth tags of images were used to supervise the training of the CNN. Not directly using the output layer of a RNN, the work in [11] utilized the Fisher vector derived from the gradient of the RNN, as the feature representation.

Although RNN is a suitable model for the sequential image annotation task for its ability to implicitly encode the dependencies among tags, it is not easy to explicitly embed some prior knowledge about the tag dependencies like se-

semantic hierarchy [26] or mutual exclusion [2] in the RNN model. To remedy this issue, the recent work of DIA [23] formulated the sequential prediction as a sampling process based on a determinantal point process (DPP) [10]. DIA encodes the class co-occurrence into the learning process, and incorporates the semantic hierarchy into the sampling process. Another important difference between DIA and the RNN-based methods is that the former explicitly embeds the negative correlations among tags *i.e.*, avoiding using semantically similar tags for the same image, while RNN-based methods typically ignore such negative correlations. The main reason is that the objective of DIA is to describe an image with a few diverse and relevant tags, while most other methods tend to predict most relevant tags.

Our proposed model D^2IA -GAN is inspired by DIA, and both are developed based on the observations of human annotations. Yet, there are several significant differences between them. The most important difference is in their objectives. DIA aims to simulate a single human annotator to use semantically distinct tags for an image, while D^2IA -GAN aims to simulate multiple human annotators simultaneously to capture the diversity among human annotators. They are also different in the training process, which will be reviewed in the Section 4. Besides, in DIA [23], ‘diverse/diversity’ refers to the semantic difference between tags in the same tag subset, to which we use the word ‘distinct/distinctiveness’ for the same meaning in this work. We use ‘diverse/diversity’ to indicate the semantic difference between multiple tag subsets for the same image.

3. Background

Weighted semantic paths. Weighted semantic paths [23] are constructed based on the semantic hierarchy and synonyms [26] among all candidate tags. To construct a weighted semantic path, we treat each tag as a node, and the synonyms are merged into one node. Then, starting from each leaf node in the semantic hierarchy, we connect its direct parent node and repeat this connection process, until the root node is achieved. All tags that are visited in this process form the weighted semantic path of the leaf tag. The weight of each tag in the semantic path is computed inversely proportional to the node layer (the layer number starts from 0 at leaf nodes) and the number of descendants of each node. As such, the weight of the tag with more specified information will be larger. A brief example of the weighted semantic paths is shown in Fig. 3. We use $SP_{\mathcal{T}}$ to denote the semantic paths of set \mathcal{T} of all candidate tags. SP_T indicates the semantic paths of the tag subset T . SP_I represents the weighted semantic paths of all ground-truth tags of image I .

Diverse and distinct tag subsets. Given an image I and its ground-truth semantic paths SP_I , a tag sub-

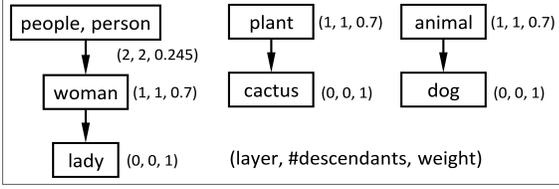


Figure 3. A brief example of the weighted semantic paths. The word in box indicates the tag. The arrow $a \rightarrow b$ tells that tag a is the semantic parent of tag b . The bracket close to each box denotes the corresponding (node layer, number of descendants, tag weight). Boxes connected by arrows construct a semantic path.

set is distinct if there are no tags being sampled from the same semantic path. An example of the distinct tag subset is shown in Fig. 3: $SP_I = \{lady \rightarrow woman \rightarrow (people, person), cactus \rightarrow plant\}$ includes 3 semantic paths with 7 tags, such as $\{lady, plant\}$ or $\{woman, cactus\}$. A tag set is diverse if it includes multiple distinct tag subsets. These subsets cover different contents of the image, due to two possible reasons, including 1) they describe different contents of the image, and 2) they describe the same content but at different semantic levels. As shown in Fig. 3, we can construct a diverse set of distinct tag subsets like $\{\{lady, cactus\}, \{plant, cat\}, \{woman, plant, animal\}\}$. Furthermore, we can construct all possible distinct tag subsets (ignoring the subset with a single tag) to obtain the complete diverse set of distinct tag subsets, referred to as S_{DD-I} . Specifically, for the subset with 2 tags, we will pick 2 paths out of 3 and sample one tag from each picked path. Then we obtain in total 16 distinct subsets. For the subset with 3 tags, we sample one tag from each semantic path, leading to 12 distinct subsets. S_{DD-I} will be used as the ground-truth to train the proposed model.

Conditional DPP. We use a conditional determinantal point process (DPP) model to measure the probability of the tag subset T , derived from the ground set \mathcal{T} given a feature \mathbf{x} of the image I . The DPP model is formulated as

$$\mathcal{P}(T|I) = \frac{\det(\mathbf{L}_T(I))}{\det(\mathbf{L}_T(I) + \mathbf{I})}, \quad (1)$$

where $\mathbf{L}_T(I) \in \mathbb{R}^{|\mathcal{T}| \times |\mathcal{T}|}$ is a positive semi-definite kernel matrix. \mathbf{I} indicates the identity matrix. For clarity, the parameters of $\mathbf{L}_T(I)$ and (1) have been omitted. The sub-matrix $\mathbf{L}_T(I) \in \mathbb{R}^{|\mathcal{T}| \times |\mathcal{T}|}$ is constructed by extracting the rows and columns corresponding to the tag indexes in T . For example, assuming $\mathbf{L}_T(I) = [a_{ij}]_{i,j=1,2,3,4}$ and $T = \{2, 4\}$, then $\mathbf{L}_T(I) = [a_{22}, a_{24}; a_{42}, a_{44}]$. $\det(\mathbf{L}_T(I))$ indicates the determinant of $\mathbf{L}_T(I)$. It encodes the negative correlations among the tags in the subset T .

Learning the kernel matrix $\mathbf{L}_T(I)$ directly is often difficult, especially when $|\mathcal{T}|$ is large. To alleviate this problem, we decompose $\mathbf{L}_T(I)$ as $\mathbf{L}_T(i, j) = v_i \phi_i^\top \phi_j v_j$, where the

scalar v_i indicates the individual score with respect to tag i , and $\mathbf{v}_T = [v_1, \dots, v_i, \dots, v_T]$. The vector $\phi_i \in \mathbb{R}^d$ corresponds to the direction of tag i , with $\|\phi_i\| = 1$, and can be used to construct the semantic similarity matrix $\mathbf{S}_T \in \mathbb{R}^{|\mathcal{T}| \times |\mathcal{T}|}$ with $\mathbf{S}_T(i, j) = \phi_i^\top \phi_j$. With this decomposition, we can learn \mathbf{v}_T and \mathbf{S}_T separately. More details of DPP can be found in [10]. In this work, \mathbf{S}_T is pre-computed as:

$$\mathbf{S}_T(i, j) = \frac{1}{2} + \frac{\langle \mathbf{t}_i, \mathbf{t}_j \rangle}{2\|\mathbf{t}_i\|_2\|\mathbf{t}_j\|_2} \in [0, 1] \quad \forall i, j \in \mathcal{T}, \quad (2)$$

where the tag representation $\mathbf{t}_i \in \mathbb{R}^{50}$ is derived from the GloVe algorithm [17]. $\langle \cdot, \cdot \rangle$ indicates the inner product of two vectors, while $\|\cdot\|_2$ denotes the ℓ_2 norm of a vector.

k-DPP sampling with weighted semantic paths. k-DPP sampling [10] is a sequential sampling process to obtain a tag subset T with at most k tags, according to the distribution (1) and the weighted semantic paths SP_T . It is denoted as $\mathcal{S}_{k\text{-DPP}, SP_T}(\mathbf{v}_T, \mathbf{S}_T)$ subsequently. Specifically, in each sampling step, the newly sampled tag will be checked whether it is from the same semantic path with any previously sampled tags. If not, it is included into the tag subset; if yes, it is abandoned and we go on sampling the next tag, until k tags are obtained. The whole sampling process is repeated multiple times to obtain different tag subsets. Then the subset with the largest tag weight summation is picked as the final output. Note that a larger weight summation indicates more semantic information. Since the tag weight is pre-defined when introducing the weighted semantic paths, it is an objective criterion to pick the subset.

4. D²IA-GAN Model

Given an image I , we aim to generate a diverse tag set including multiple distinct tag subsets relevant to the image content, as well as an ensemble tag subset of these distinct subsets, which could provide a comprehensive description of I . These tags are sampled from a generative model conditioned on the image, and we use a conditional GAN (CGAN) [16, 30, 1] to train it, with the generator part \mathcal{G} being our model and a discriminator \mathcal{D} , as shown in Fig. 2. Specifically, conditioned on I , \mathcal{G} projects one noise vector \mathbf{z} to one distinct tag subset T , and uses different noise vectors to ensure diverse/different tag subsets. \mathcal{D} serves as an adversary of \mathcal{G} , aiming to distinguish the generated tag subsets using \mathcal{G} from the ground-truth ones S_{DD-I} .

4.1. Generator

The tag subset $T \subset \mathcal{T} = \{1, 2, \dots, m\}$ with $|T| \leq k$ can be generated from the generator $G_\theta(I, \mathbf{z})$, according to the input image I and a noise vector \mathbf{z} , as follows:

$$G_\theta(I, \mathbf{z}; \mathbf{S}_T, SP_T, k) \sim \mathcal{S}_{k\text{-DPP}, SP_T}(\sqrt{\mathbf{q}_T(I, \mathbf{z})}, \mathbf{S}_T). \quad (3)$$

The above generator is a composite function with two parts. The **inner** part $\mathbf{q}_{\mathcal{T}}(I, \mathbf{z}) = \sigma(\mathbf{W}^{\top}[f_{\mathcal{G}}(I); \mathbf{z}] + \mathbf{b}_{\mathcal{G}}) \in [0, 1]^{|T|}$ is a CNN based soft classifier. $f_{\mathcal{G}}(I)$ represents the output vector of the fully-connected layer of a CNN model, and $[\mathbf{a}_1; \mathbf{a}_2]$ denotes the concatenation of two vectors \mathbf{a}_1 and \mathbf{a}_2 . $\sigma(\mathbf{a}) = \frac{1}{1+\exp(-\mathbf{a})}$ is the sigmoid function. $\sqrt{\mathbf{a}}$ indicates the element-wise square root of vector \mathbf{a} . The parameter matrix $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_i, \dots, \mathbf{w}_m] \in \mathbb{R}^{m \times d}$ and the bias parameter $\mathbf{b}_{\mathcal{G}} \in \mathbb{R}^d$ map the feature vector $[f_{\mathcal{G}}(I); \mathbf{z}] \in \mathbb{R}^d$ to the logit vector. The trainable parameter θ includes \mathbf{W} , $\mathbf{b}_{\mathcal{G}}$ and the parameters of $f_{\mathcal{G}}$. The noise vector \mathbf{z} is sampled from the uniform distribution $\mathcal{U}[-1, 1]$. The **outer** part $\mathcal{S}_{k\text{-DPP}, SP_{\mathcal{T}}}(\sqrt{\mathbf{q}_{\mathcal{T}}(I, \mathbf{z})}, \mathbf{S}_{\mathcal{T}})$ is the k -DPP sampling with weighted semantic paths $SP_{\mathcal{T}}$ (see Section 3). Using $\sqrt{\mathbf{q}_{\mathcal{T}}(I, \mathbf{z})}$ as the quality term and utilizing the pre-defined similarity matrix $\mathbf{S}_{\mathcal{T}}$, then a conditional DPP model can be constructed as described in Section 3.

4.2. Discriminator

$\mathcal{D}_{\eta}(I, T)$ evaluates the relevance of image I and tag subset T : it outputs a value in $[0, 1]$, with 1 meaning the highest relevance and 0 being the least relevant. Specifically, \mathcal{D}_{η} is constructed as follows: first, as described in Section 3, each tag $i \in T$ is represented by a vector $\mathbf{t}_i \in \mathbb{R}^{50}$ derived from the GloVe algorithm [17]. Then, we formulate $\mathcal{D}_{\eta}(I, T)$ as

$$\mathcal{D}_{\eta}(I, T) = \frac{1}{|T|} \sum_{i \in T} \sigma(\mathbf{w}_{\mathcal{D}}^{\top}[f_{\mathcal{D}}(I); \mathbf{t}_i] + b_{\mathcal{D}}), \quad (4)$$

where $f_{\mathcal{D}}(I)$ denotes the output vector of the fully-connected layer of a CNN model (different from that used in the generator). η includes $\mathbf{w}_{\mathcal{D}} \in \mathbb{R}^{|f_{\mathcal{D}}(I)|+50}$, $b_{\mathcal{D}} \in \mathbb{R}$ and the parameters of $f_{\mathcal{D}}(I)$ in the CNN model.

4.3. Conditional GAN

Following the general training procedure, we learn D²IA-GAN by iterating two steps until convergence: (1) fixing the discriminator \mathcal{D}_{η} and optimizing the generator \mathcal{G}_{θ} using (5), as shown in Section 4.3.1; (2) fixing \mathcal{G}_{θ} and optimizing \mathcal{D}_{η} using (8), as shown in Section 4.3.2.

4.3.1 Optimizing \mathcal{G}_{θ}

Given \mathcal{D}_{η} , we learn \mathcal{G}_{θ} by

$$\min_{\theta} \mathbb{E}_{\mathbf{z} \sim \mathcal{U}[-1, 1]} \left[\log \left(1 - \mathcal{D}_{\eta}(I, \mathcal{G}_{\theta}(I, \mathbf{z})) \right) \right]. \quad (5)$$

For clarity, we only show the case with one training image I in the above formulation. Due to the discrete sampling process $\mathcal{S}(\mathbf{v}_{\mathcal{T}}, \mathbf{S}_{\mathcal{T}})$ in $\mathcal{G}_{\theta}(I, \mathbf{z})$, we cannot optimize (5) using any existing continuous optimization algorithm. To address this issue, we view the sequential generation of tags as controlled by a continuous policy function, which weighs different choices of the next tag based on the image and tags already generated. As such, we can use the

policy gradient (PG) algorithm in reinforcement learning for its optimization. Given a sampled tag subset $T_{\mathcal{G}}$ from $\mathcal{S}(\mathbf{v}_{\mathcal{T}}, \mathbf{S}_{\mathcal{T}})$, the original objective function of (5) is approximated by a continuous function. Specifically, we denote $T_{\mathcal{G}} = \{y_{[1]}, y_{[2]}, \dots, y_{[k]}\}$, where $[i]$ indicates the sampling order, and its subset $T_{\mathcal{G}-i} = \{y_{[1]}, \dots, y_{[i]}\}$, $i \leq k$ includes the first i tags in $T_{\mathcal{G}}$. Then, with an instantiated \mathbf{z} sampled from $[-1, 1]$, the approximated function is formulated as

$$\mathcal{J}_{\theta}(T_{\mathcal{G}}) = \sum_{i=1}^k \mathcal{R}(I, T_{\mathcal{G}-i}) \log \left(\prod_{t_1 \in T_{\mathcal{G}-i}} q_{t_1}^1 \prod_{t_2 \in \mathcal{T} \setminus T_{\mathcal{G}-i}} q_{t_2}^0 \right), \quad (6)$$

where $\mathcal{T} \setminus T_{\mathcal{G}-i}$ denotes the relative complement of $T_{\mathcal{G}-i}$ with respect to \mathcal{T} . $q_t^1 = \sigma(\mathbf{w}_t^{\top}[f_{\mathcal{G}}(I); \mathbf{z}] + b_{\mathcal{G}}(t))$ indicates the posterior probability, and $q_t^0 = 1 - q_t^1$. The reward function $\mathcal{R}(I, T_{\mathcal{G}})$ encourages the content of I and the tags $T_{\mathcal{G}}$ to be consistent, and is defined as

$$\mathcal{R}(I, T_{\mathcal{G}}) = -\log(1 - \mathcal{D}_{\eta}(I, T_{\mathcal{G}})). \quad (7)$$

Compared to a full PG objective function, in (6) we have replaced the *return* with the *immediate reward* $\mathcal{R}(I, T_{\mathcal{G}})$, and the *policy* probability with the decomposed likelihood $\prod_{t_1 \in T_{\mathcal{G}-i}} q_{t_1}^1 \prod_{t_2 \in \mathcal{T} \setminus T_{\mathcal{G}-i}} q_{t_2}^0$. Consequently, it is easy to compute the gradient $\frac{\partial \mathcal{J}_{\theta}(T_{\mathcal{G}})}{\partial \theta}$, which will be used in the stochastic gradient ascent algorithm and back-propagation [19] to update θ .

When generating $T_{\mathcal{G}}$ during training, we repeat the sampling process multiple times to obtain different subsets. Then, as the ground-truth set S_{DD-I} for each training image is available, the semantic F_{1-sp} score (see Section 5) for each generated subset can be computed, and the one with the largest F_{1-sp} score will be used to update parameters. This process encourages the model to generate tag subsets more consistent with the evaluation metric.

4.3.2 Optimizing \mathcal{D}_{η}

Utilizing the generated tag subset $T_{\mathcal{G}}$ from the fixed generator $\mathcal{G}_{\theta}(I, \mathbf{z})$, we learn \mathcal{D}_{η} by

$$\max_{\eta} \frac{1}{|S_{DD-I}|} \sum_{T \in S_{DD-I}} \left[\beta \log \mathcal{D}_{\eta}(I, T) - (1 - \beta) \cdot \left(\mathcal{D}_{\eta}(I, T) - F_{1-sp}(I, T) \right)^2 \right] + \beta \log(1 - \mathcal{D}_{\eta}(I, T_{\mathcal{G}})) - (1 - \beta) \left(\mathcal{D}_{\eta}(I, T_{\mathcal{G}}) - F_{1-sp}(I, T_{\mathcal{G}}) \right)^2, \quad (8)$$

where semantic score $F_{1-sp}(I, T)$ measures the relevance between the tag subset T and the content of I . If we set the trade-off parameter $\beta = 1$, then (8) is equivalent to the objective used in the standard GAN model. For $\beta \in (0, 1)$, (8) also encourages the updated \mathcal{D}_{η} to be close to the semantic score $F_{1-sp}(I, T)$. We can then compute the gradient of (8) with respect to η , and use the stochastic gradient ascent algorithm and back-propagation [19] to update η .

5. Experiments

5.1. Experimental Settings

Datasets. We adopt two benchmark datasets, ESP Game [20] and IAPRTC-12 [8] for evaluation. One important reason for choosing these two datasets is that they have complete weighted semantic paths of all candidate tags SP_T , the ground-truth weighted semantic paths of each image SP_I , the image features and the trained DIA model, provided by the authors of [23] and available on GitHub¹. Since the weighted semantic paths are important to our method, these two datasets facilitate the evaluation. Specifically, in ESP Game, there are 18689 train images, 2081 test images, 268 candidate classes, 106 semantic paths corresponding to all candidate tags, and the feature dimension is 597; in IAPRTC-12, there are 17495 train images, 1957 test images, 291 candidate classes, 139 semantic paths of all candidate tags, and the feature dimension is 536.

Model training. We firstly fix the CNN models in both \mathcal{G}_θ and \mathcal{D}_η as the VGG-F model² pre-trained on ImageNet [3]. Then we initialize the columns of the fully-connected parameter matrix \mathbf{W} (see Eq. (3)) that corresponds to the image feature $f_G(I)$ using the trained DIA model, while the columns corresponding to the noise vector \mathbf{z} and the bias parameter \mathbf{b}_G are randomly initialized. We pre-train \mathcal{D}_η by setting $\beta = 0$ in Eq. (8), *i.e.*, only using the F_{1-sp} scores of ground-truth subsets S_{DD-I} and the fake subsets generated by the initialized \mathcal{G}_θ with \mathbf{z} being the zero vector. The corresponding pre-training parameters are: batch size = 256, epochs = 20, learning rate = 1, ℓ_2 weight decay = 0.0001. With the initialized \mathcal{G}_θ and the pre-trained \mathcal{D}_η , we fine-tune the D²IA-GAN model using the following parameters: batch size = 256, epochs = 50, the learning rates of \mathbf{W} and η are respectively set to 0.0001 and 0.00005, both learning rates are decayed by 0.1 in every 10 epochs, ℓ_2 weight decay = 0.0001, and $\beta = 0.5$. Besides, if there are a few long paths (*i.e.*, many tags in a semantic path) in SP_I , the number of subsets in SP_I , *i.e.*, $|S_{DD-I}|$, could be very large. In ESP Game and IAPRTC-12, the largest $|S_{DD-I}|$ is up to 4000, though the $|S_{DD-I}|$ values for most images are smaller than 30. If $|S_{DD-I}|$ is too large, the training of the discriminator \mathcal{D}_η (see Eq. (8)) will be slow. Thus, we set an upper bound 10 for $|S_{DD-I}|$ in training, if $|S_{DD-I}| > 10$, then we randomly choose 10 subsets from S_{DD-I} to update \mathcal{D}_η . The implementation adopts Tensorflow 1.2.0 and Python 2.7.

Evaluation metrics. To evaluate the distinctiveness and relevance of the predicted tag subset, three semantic metrics, including semantic precision, recall and F1, are proposed in [23], according to the weighted semantic paths.

They are denoted as P_{sp} , R_{sp} and F_{1-sp} , respectively. Specifically, given a predicted subset T , the corresponding semantic paths SP_T and the ground-truth semantic paths SP_I , P_{sp} computes the proportion of the true semantic paths in SP_T , and R_{sp} computes the proportion of the true semantic paths in SP_I that are also included in SP_T , and $F_{1-sp} = 2(P_{sp} \cdot R_{sp}) / (P_{sp} + R_{sp})$. The tag weights in each path are also considered when computing the proportions. Please refer to [23] for the detailed definitions.

Comparisons. We compare with two state-of-the-art image annotation methods, including ML-MG³ [26] and DIA⁴ [23]. The reason we compare with them is that both they and our method utilize the semantic hierarchy and the weighted semantic paths, though with different usages. We also compare with another state-of-the-art multi-label learning method, called LEML⁵ [31], which doesn't utilize the semantic hierarchy. Since both ML-MG and LEML do not consider the semantic distinctiveness among tags, their predicted tag subsets are likely to include semantic redundancies. As reported in [23], the evaluation scores using the semantic metrics (*i.e.*, P_{sp} , R_{sp} and F_{1-sp}) of ML-MG and LEML's predictions are much lower than DIA. Hence it is not relevant to compare with the original results of ML-MG and LEML. Instead, we combine the predictions of ML-MG and LEML with the DPP-sampling that is also used in DIA and our method. Specifically, the square root of posterior probabilities with respect to all candidate tags produced by ML-MG are used as the quality vector (see Section 3); as there are negative scores in the predictions of LEML, we normalize all predicted scores to $[0, 1]$ to obtain the posterior probabilities. Then combining with the similarity matrix \mathbf{S} , two DPP distributions are constructed to sample distinct tag subsets. The obtained results denoted as MLMG-DPP and LEML-DPP, respectively.

5.2. Quantitative Results

As all compared methods (MLMG-DPP, LEML-DPP and DIA) and the proposed D²IA-GAN sample from DPP distributions to generate tag subsets, we can generate multiple tag subsets using each method for each image. Specifically, MLMG-DPP and DIA generate 10 random tag subsets for each image. The weight of each tag subset is computed by summing the weights of all tags in the subset. Then, we construct two outputs, including: the *single subset*, which picks the subset with the largest weight from these 10 subsets; and the *ensemble subset*, which merges 5 tag subsets with top-5 largest weights among 10 subsets into one unique tag subset. The evaluation of the single subset reflects the performance of distinctiveness of the compared methods. The evaluation of the ensemble subset measures the per-

¹Downloaded from <https://github.com/wubaoyuan/DIA>

²Downloaded from <http://www.vlfeat.org/matconvnet/pretrained/>

³Downloaded from <https://sites.google.com/site/baoyuanwu2015/home>

⁴Downloaded from <https://github.com/wubaoyuan/DIA>

⁵Downloaded from <http://www.cs.utexas.edu/~rofyuy/>

evaluation target	metric→ method↓	3 tags			5 tags		
		P_{sp}	R_{sp}	F_{1-sp}	P_{sp}	R_{sp}	F_{1-sp}
single subset	LEML-DPP [31]	34.64	25.21	27.76	29.24	35.05	30.29
	MLMG-DPP [26]	37.18	27.71	30.05	33.85	38.91	34.30
	DIA [23]	41.44	31.00	33.61	34.99	40.92	35.78
	D ² IA-GAN	42.96	32.34	34.93	35.04	41.50	36.06
ensemble subset	LEML-DPP [31]	34.62	38.09	34.32	29.04	46.61	34.02
	MLMG-DPP [26]	30.44	34.88	30.70	28.99	43.46	33.05
	DIA [23]	35.73	33.53	32.39	32.62	40.86	34.31
	D ² IA-GAN	36.73	42.44	36.71	31.28	48.74	35.82

Table 1. Results (%) evaluated by semantic metrics on ESP Game. The higher value indicates the better performance, and the best result in each column is highlighted in bold.

evaluation target	metric→ method↓	3 tags			5 tags		
		P_{sp}	R_{sp}	F_{1-sp}	P_{sp}	R_{sp}	F_{1-sp}
single subset	LEML-DPP [31]	41.42	24.39	29.00	37.06	32.86	32.98
	MLMG-DPP [26]	40.93	24.29	28.61	37.06	33.68	33.29
	DIA [23]	42.65	25.07	29.87	37.83	34.62	34.11
	D ² IA-GAN	43.57	26.22	31.04	37.31	35.35	34.41
ensemble subset	LEML-DPP [31]	35.22	32.75	31.86	32.28	39.89	33.74
	MLMG-DPP [26]	33.71	32.00	30.64	31.91	40.11	33.49
	DIA [23]	35.73	33.53	32.39	32.62	40.86	34.31
	D ² IA-GAN	35.49	39.06	34.44	32.50	44.98	35.34

Table 2. Results (%) evaluated by semantic metrics on IAPRTC-12. The higher value indicates the better performance, and the best result in each column is highlighted in bold.

formance of both diversity and distinctiveness. Larger distinctiveness of the ensemble subset indicates higher diversity among the consisting subsets of this ensemble subset. Moreover, we present two cases by limiting the size of each tag subset to 3 and 5, respectively.

The quantitative results on ESP Game are shown in Table 1. For evaluating single subsets, D²IA-GAN shows the best performance evaluated by all metrics for both 3 and 5 tags, while MLMG-DPP and LEML-DPP perform worst in all cases. The reason is that the learning of ML-MG/LEML and the DPP sampling are independent. For ML-MG, it enforces the ancestor tags to be ranked before their descendant tags, while the distinctiveness is not considered. There is much semantic redundancy in top-k tags of ML-MG, which is likely to include fewer semantic paths than those of DIA and D²IA-GAN. Hence, although DPP sampling can produce a distinct tag subset from top-k candidate tags, it covers fewer semantic concepts (remember that one semantic path represents one semantic concept) than DIA and D²IA-GAN. For LEML, it treats each tag equally when training, while totally ignoring the semantic distinctiveness. It is not surprising that LEML-DPP also covers fewer semantic concepts than DIA and D²IA-GAN. In contrast, both DIA and D²IA-GAN take into account the semantic distinctiveness in learning. However, there are several significant differences between their training processes. First, DPP sampling is independent with the model training in DIA, while the generated subset by DPP sampling is used to update the model parameters in D²IA-GAN. Second, DIA learns from the ground-truth complete tag list, and the semantic distinctiveness is indirectly embedded into the learning process through the similarity matrix S . In

contrast, D²IA-GAN learns from the ground-truth distinct tag subsets. Last, the model training of DIA is independent of the evaluation metric F_{1-sp} , which plays the important role in the training of D²IA-GAN. These differences are the causes that D²IA-GAN produces more semantically distinct tag subsets than DIA. Specifically, in the case of 3 tags, the relative improvements of D²IA-GAN over DIA are 3.67%, 4.32%, 3.93% at P_{sp} , R_{sp} and F_{1-sp} , respectively; while being 0.14%, 3.86% and 0.78% in the case of 5 tags. In addition, the improvement decreases as the size limit of tag subset increases. The reason is that D²IA-GAN may include more irrelevant tags, as the random noise combined with the image feature not only brings in diversity, but also uncertainty. Note that due to the randomness of sampling, the results of single subset by DIA presented here are slightly different with those reported in [23].

In terms of the evaluation of ensemble subsets, the improvements of D²IA-GAN over three compared methods are more significant. This is because all three compared methods sample multiple tag subsets from a fixed DPP distribution, while D²IA-GAN generates multiple tag subsets from different DPP distributions with random perturbations. As such, the diversity among tag subsets generated by D²IA-GAN is expected to be higher than those corresponding to three compared methods. Subsequently, the ensemble subset of D²IA-GAN is likely to cover more relevant semantic paths than those of other methods. It is supported by the comparison through the evaluation by R_{sp} : the relative improvement of D²IA-GAN over DIA is 26.57% in the case of 3 tags, while 19.29% in the case of 5 tags. It is encouraging that the P_{sp} scores of D²IA-GAN are also comparable with those of DIA. It demonstrates that training using GAN reduces the likelihood to include irrelevant semantic paths due to the uncertainty of the noise vector z , because GAN encourages the generated tag subsets to be close to the ground-truth diverse and distinct tag subsets. Specifically, in the case of 3 tags, the relative improvements of D²IA-GAN over DIA are 2.80%, 26.57%, 11.77% at P_{sp} , R_{sp} and F_{1-sp} , respectively; the corresponding improvements are -4.11%, 19.29%, 4.40% in the case of 5 tags.

The results on IAPRTC-12 are summarized in Table 2. In the case of single subset with 3 tags, the relative improvements of D²IA-GAN over DIA are 2.16%, 4.59%, 3.92% at P_{sp} , R_{sp} and F_{1-sp} , respectively. In the case of single subset with 5 tags, the corresponding improvements are -1.37%, 2.11%, 0.88%. In the case of ensemble subset and 3 tags, the corresponding improvements are -0.67%, 16.49%, 6.33%. In the case of ensemble subset and 5 tags, the corresponding improvements are -0.37%, 10.08%, 3.0%. The comparisons on above two benchmark datasets verify that D²IA-GAN produces more semantically relevant, yet diverse and distinct tag subsets than the compared MLMG-DPP, LEML-DPP and DIA

methods. Some qualitative results will be presented in the **supplementary material**.

5.3. Subject Study

Since the diversity and distinctiveness are subjective concepts, we also conduct human subject studies to compare the results of DIA and D²-GAN on these two criterion. Specifically, for each test image, we run DIA 10 times to obtain 10 tag subsets, and then the set including 3 subsets with the largest weights are picked as the final output. For D²-GAN, we firstly generate 10 random noise vectors \mathbf{z} . With each noise vector, we conduct the DPP sampling in \mathcal{G}_θ for 10 times to obtain 10 subsets, out of which we pick that with the largest weight as the tag subset corresponding to this noise vector. Then, from the obtained 10 subsets, we again pick 3 subsets with the largest weights to form the output set of D²-GAN. For each test image, we present these two sets of tag subsets with the corresponding image to 5 human evaluators. The only instruction to evaluators is to determine “which set describes this image more comprehensively”. Besides, we notice that if two tag sets are very similar, or if they both are irrelevant to the image contents, human evaluators may pick one randomly. To reduce such randomness, we filter the test images using the following criterion: firstly we combine the subsets in each set to an ensemble subset; if the F_{1-sp} scores of both ensemble subsets are larger than 0.2, and the gap between this two scores is larger than 0.15, then this image is used in subject studies. Finally, the numbers of test images used in subject studies are: ESP Game, 375 in the case of 3 tags, and 324 in the case of 5 tags; IAPRTC-12, 342 in the case of 3 tags, and 306 in the case of 5 tags. We also present the comparison results using F_{1-sp} to evaluate the compared two ensemble subsets. The consistency between the F_{1-sp} evaluation and the human evaluation is computed. Subject study results on ESP Game are summarized in Table 3. With human evaluation, D²IA-GAN is judged better at $\frac{240}{375} = 64\%$ of all evaluated images over DIA in the case of 3 tags, and $\frac{204}{324} = 62.96\%$ in the case of 5 tags. With F_{1-sp} evaluation, D²IA-GAN outperforms DIA at $\frac{250}{375} = 66.67\%$ in the case of 3 tags, and $\frac{212}{324} = 65.43\%$ in the case of 5 tags. Both evaluation results suggest the improvement of D²IA-GAN over DIA. Besides, the results of these two evaluations are consistent (*i.e.*, their decisions of which set is better are same) at $\frac{239}{375} = 63.73\%$ of all evaluated images in the case of 3 tags, while $\frac{222}{324} = 68.52\%$ in the case of 5 tags. It demonstrates that the evaluation using F_{1-sp} is relatively reliable. The same trend is also observed for the results obtained on the IAPRTC-12 dataset (see Table 4).

Moreover, in the **supplementary material**, we will present a detailed analysis about human annotations conducted on partial images of IAPRTC-12. It not only shows that D²IA-GAN produces more human-like tags than DIA, but also discusses the difference between D²IA-GAN and

# tags →	3 tags			5 tags		
	DIA wins	D ² IA-GAN wins	total	DIA wins	D ² IA-GAN wins	total
human evaluation	135	240	375	120	204	324
F_{1-sp}	125	250	375	112	212	324
consistency	62	177	63.73%	65	157	68.52%

Table 3. Subject study results on ESP Game. Note that the entry ‘62’ corresponding to the row ‘consistency’ and the column ‘DIA wins’ indicates that both human evaluation and F_{1-sp} evaluation decide that the predicted tags of DIA are better than those of D²IA-GAN at 62 images. Similarly, human evaluation and F_{1-sp} evaluation have the same decision that the results of D²IA-GAN are better than those of DIA at 177 images. Hence, two evaluations have the same decision (*i.e.*, consistent) on $62 + 177 = 239$ images, and the consistency rate among all evaluated images are $239/375 = 63.73\%$.

# tags →	3 tags			5 tags		
	DIA wins	D ² IA-GAN wins	total	DIA wins	D ² IA-GAN wins	total
human evaluation	129	213	342	123	183	306
F_{1-sp}	141	201	342	123	183	306
consistency	82	154	69.01%	58	118	57.52%

Table 4. Subject study results on IAPRTC-12.

human annotators, as well as how to shrink that difference.

6. Conclusion

In this work, we have proposed a new image annotation method, called *diverse and distinct image annotation* (D²IA), to simulate the diversity and distinctiveness of the tags generated by human annotators. D²IA is formulated as a sequential generative model, in which the image feature is firstly incorporated into a determinantal point process (DPP) model that also encodes the weighted semantic paths, from which a sequence of distinct tags are generated by sampling. The diversity among the generated multiple tag subsets is ensured by sampling the DPP model with random noise perturbations to the image feature. In addition, we adopt the generative adversarial network (GAN) model to train the generative model D²IA, and employ the policy gradient algorithm to handle the training difficulty due to the discrete DPP sampling in D²IA. Experimental results and human subject studies on two benchmark datasets demonstrate that the relevant, yet diverse and distinct tag subsets generated by the proposed method D²IA-GAN can provide more comprehensive descriptions of the image contents than those generated by state-of-the-art methods.

Acknowledgements: This work is supported by Tencent AI Lab. The participation of Bernard Ghanem is supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research. The participation of Siwei Lyu is partially supported by National Science Foundation National Robotics Initiative (NRI) Grant (IIS-1537257) and National Science Foundation of China Project Number 61771341.

References

- [1] Y. Bai, Y. Zhang, M. Ding, and B. Ghanem. Finding tiny faces in the wild with generative adversarial network. In *CVPR*. IEEE, 2018.
- [2] X. Chen, X.-T. Yuan, Q. Chen, S. Yan, and T.-S. Chua. Multi-label visual classification with label exclusive context. In *ICCV*, pages 834–841, 2011.
- [3] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. IEEE, 2009.
- [4] C. Gong, D. Tao, W. Liu, L. Liu, and J. Yang. Label propagation via teaching-to-learn and learning-to-teach. *IEEE transactions on neural networks and learning systems*, 28(6):1452–1465, 2017.
- [5] C. Gong, D. Tao, J. Yang, and W. Liu. Teaching-to-learn and learning-to-teach for multi-label propagation. In *AAAI*, pages 1610–1616, 2016.
- [6] Y. Gong, Y. Jia, T. Leung, A. Toshev, and S. Ioffe. Deep convolutional ranking for multilabel image annotation. *arXiv preprint arXiv:1312.4894*, 2013.
- [7] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, pages 2672–2680, 2014.
- [8] M. Grubinger, P. Clough, H. Müller, and T. Deselaers. The iapr tc-12 benchmark: A new evaluation resource for visual information systems. In *International Workshop OntoImage*, pages 13–23, 2006.
- [9] J. Jin and H. Nakayama. Annotation order matters: Recurrent image annotator for arbitrary length image tagging. In *ICPR*, pages 2452–2457. IEEE, 2016.
- [10] A. Kulesza and B. Taskar. Determinantal point processes for machine learning. *arXiv preprint arXiv:1207.6083*, 2012.
- [11] G. Lev, G. Sadeh, B. Klein, and L. Wolf. Rnn fisher vectors for action recognition and image annotation. In *ECCV*, pages 833–850. Springer, 2016.
- [12] Y. Li, Y. Song, and J. Luo. Improving pairwise ranking for multi-label image classification. In *CVPR*, 2017.
- [13] Y. Li, B. Wu, B. Ghanem, Y. Zhao, H. Yao, and Q. Ji. Facial action unit recognition under incomplete data based on multi-label learning with missing labels. *Pattern Recognition*, 60:890–900, 2016.
- [14] F. Liu, T. Xiang, T. M. Hospedales, W. Yang, and C. Sun. Semantic regularisation for recurrent image annotation. In *CVPR*, 2017.
- [15] W. Liu, J. He, and S.-F. Chang. Large graph construction for scalable semi-supervised learning. In *ICML*, 2010.
- [16] M. Mirza and S. Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [17] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *EMNLP*, volume 14, pages 1532–43, 2014.
- [18] G.-J. Qi, W. Liu, C. Aggarwal, and T. Huang. Joint inter-modal and intramodal label transfers for extremely rare or unseen classes. *IEEE transactions on pattern analysis and machine intelligence*, 39(7):1360–1373, 2017.
- [19] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *Cognitive modeling*, 5(3):1, 1988.
- [20] L. Von Ahn and L. Dabbish. Labeling images with a computer game. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 319–326. ACM, 2004.
- [21] J. Wang, Y. Yang, J. Mao, Z. Huang, C. Huang, and W. Xu. Cnn-rnn: A unified framework for multi-label image classification. In *CVPR*, pages 2285–2294, 2016.
- [22] Z. Wang, T. Chen, G. Li, R. Xu, and L. Lin. Multi-label image recognition by recurrently discovering attentional regions. In *CVPR*, pages 464–472, 2017.
- [23] B. Wu, F. Jia, W. Liu, and B. Ghanem. Diverse image annotation. In *CVPR*, 2017.
- [24] B. Wu, F. Jia, W. Liu, B. Ghanem, and S. Lyu. Multi-label learning with missing labels using mixed dependency graphs. *International Journal of Computer Vision*, 2018.
- [25] B. Wu, Z. Liu, S. Wang, B.-G. Hu, and Q. Ji. Multi-label learning with missing labels. In *ICPR*, 2014.
- [26] B. Wu, S. Lyu, and B. Ghanem. MI-mg: Multi-label learning with missing labels using a mixed graph. In *ICCV*, pages 4157–4165, 2015.
- [27] B. Wu, S. Lyu, and B. Ghanem. Constrained submodular minimization for missing labels and class imbalance in multi-label learning. In *AAAI*, pages 2229–2236, 2016.
- [28] B. Wu, S. Lyu, B.-G. Hu, and Q. Ji. Multi-label learning with missing labels for image annotation and facial action unit recognition. *Pattern Recognition*, 48(7):2279–2289, 2015.
- [29] P. Xie, R. Salakhutdinov, L. Mou, and E. P. Xing. Deep determinantal point process for large-scale multi-label classification. In *CVPR*, pages 473–482, 2017.
- [30] W. Xiong, W. Luo, L. Ma, W. Liu, and J. Luo. Learning to generate time-lapse videos using multi-stage dynamic generative adversarial networks. In *CVPR*, 2018.
- [31] H.-F. Yu, P. Jain, P. Kar, and I. Dhillon. Large-scale multi-label learning with missing labels. In *ICML*, pages 593–601, 2014.
- [32] D. Zhang, M. M. Islam, and G. Lu. A review on automatic image annotation techniques. *Pattern Recognition*, 45(1):346–362, 2012.
- [33] F. Zhu, H. Li, W. Ouyang, N. Yu, and X. Wang. Learning spatial regularization with image-level supervisions for multi-label image classification. In *CVPR*, 2017.